Dynamic Probabilistic Relational Models

Sumit Sanghai   Pedro Domingos   Daniel Weld
Department of Computer Science and Engineering
University of Washington
Seattle, WA 98195-2350

Presented at OSU, for CS539: Probabilistic Relational Models,
Tuesday November 25th, 2003
by Matthieu Labbe.

Abusive notations in this presentation are mine.
The article is quoted without quotation mark.

How to represent/model uncertain <u>*sequential*</u> phenomena?

"State of the Art":
Dynamic Bayesian Networks (DBNs)

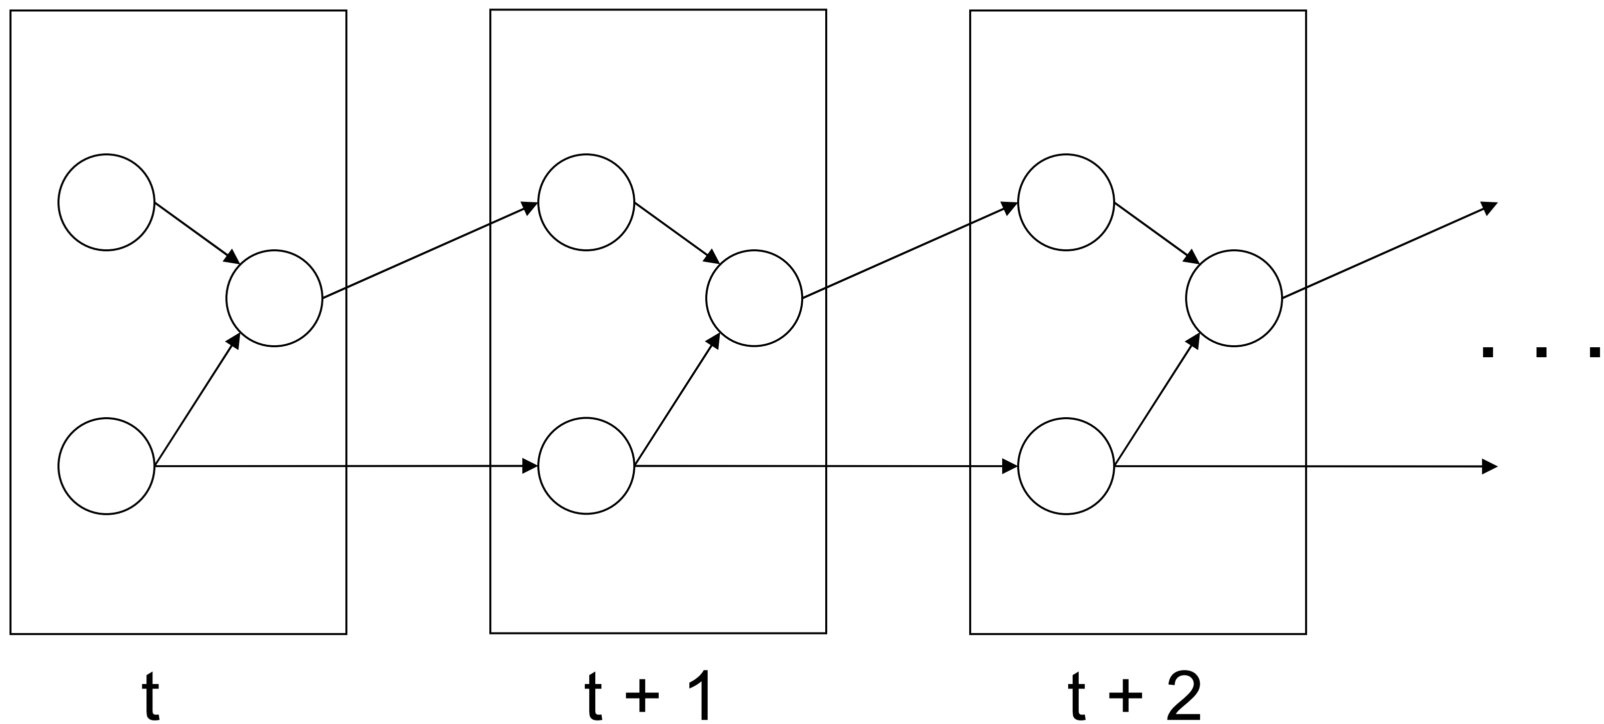

Limitations of DBNs

How to represent:

- Classes of objects and multiple instances of a class
- Multiple kinds of relations
- Relations evolving over time

Example: Early fault detection in manufacturing

Complex and diverse relations evolving over the manufacturing process.

The idea

Sequential

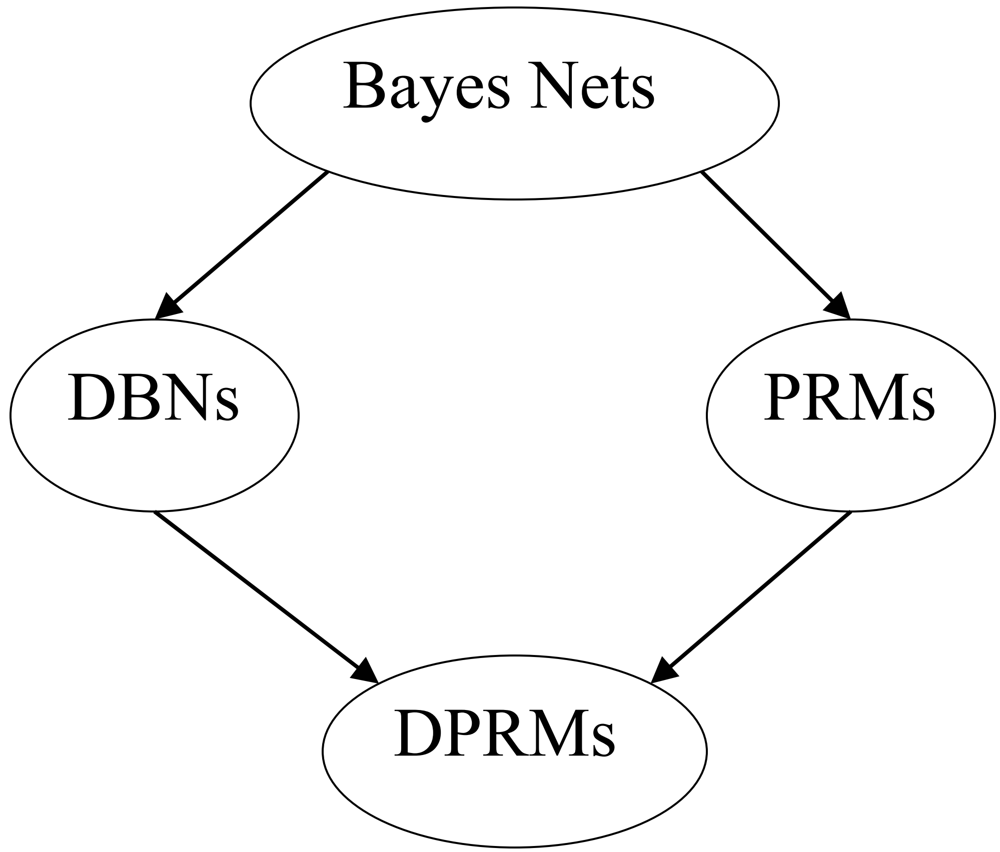

Relational

Sequential & Relational

The math

BN = (Z_i , Pa(Z_i) , P(Z_i | Pa(Z_i))) Must be a DAG.

DBN = ($Z_{i,t}$, Pa($Z_{i,t}$) , P($Z_{i,t}$ | Pa($Z_{i,t}$))) Must be a DAG.

Assuming the system is first-order Markovian and stationary:

DBN = (BN_0 , $BN_{->}$), where:

BN_0 is a BN over $Z_{i,0}$
$BN_{->}$ is a BN over $Z_{i,t+1}$ U Pa($Z_{i,t+1}$)

Using slot chain conventions:

PRM = (C_i, A(C_i), R(C_i), Pa($C_i.A_j$), P($C_i.A_j$ | Pa($C_i.A_j$)))

Key observation

DBNs as a special case of PRMs

PRM = (C, Z_i, previous, Pa(C.Z_j), P(C.Z_j | Pa(C.Z_j)))

First-order Markovian

⇔

Slot chains contain at most one “previous”

DPRMs

Definition 2 A *two-time-slice PRM (2TPRM)* for a relational schema S is defined as follows. For each class C and each propositional attribute A in A(C), we have:

- A *set of parents* $Pa(C.A) = \{Pa_1, Pa_2, ..., Pa_l\}$, where each Pa_i has the form C.B or f(C.τ.B), where τ is a slot chain containing the attribute *previous* at most once, and f() is an aggregation function.
- A *conditional probability model* for P(C.A|Pa(C.A)).

Definition 3 A *dynamic probabilistic relational model (DPRM)* for a relational schema S is a pair $(M_0, M_{->})$, where M_0 is a PRM over I_0, representing the distribution P_0 the initial instantiation of S, and $M_{->}$ is a 2TPRM representing the transition distribution $P(I_t \mid I_{t-1})$ connecting successive instantiations of S.

Inference in DPRMs

A DPRM unrolls into a DBN just as a PRM unrolls into a Bayes Net.

Inference in DBNs can be done with particle filtering, but this didn’t scale with DPRMs state space.

Rao-Blackwellisation

Rao-Blackwellisation is a technique improving particle filtering by analytically marginalizing out some of the variables.

Let's divide the variables Z_i into observed, named Y_i and unobserved, named X_i

If the variables X_i can be divided in U_i and V_i such that $P(V_i | U_i, Y_1, . . . , Y_n)$ can be computed analytically and efficiently, then we only need to sample from U_i.

Restricting assumptions

1. Relational attributes with unknown values do not appear anywhere in the DPRM as parents of unobserved attributes, or in their slot chains.
2. Each reference slot can be occupied by at most one object.

With those restrictions, we can “Rao-Blackwellise out” all relational attributes.

Results

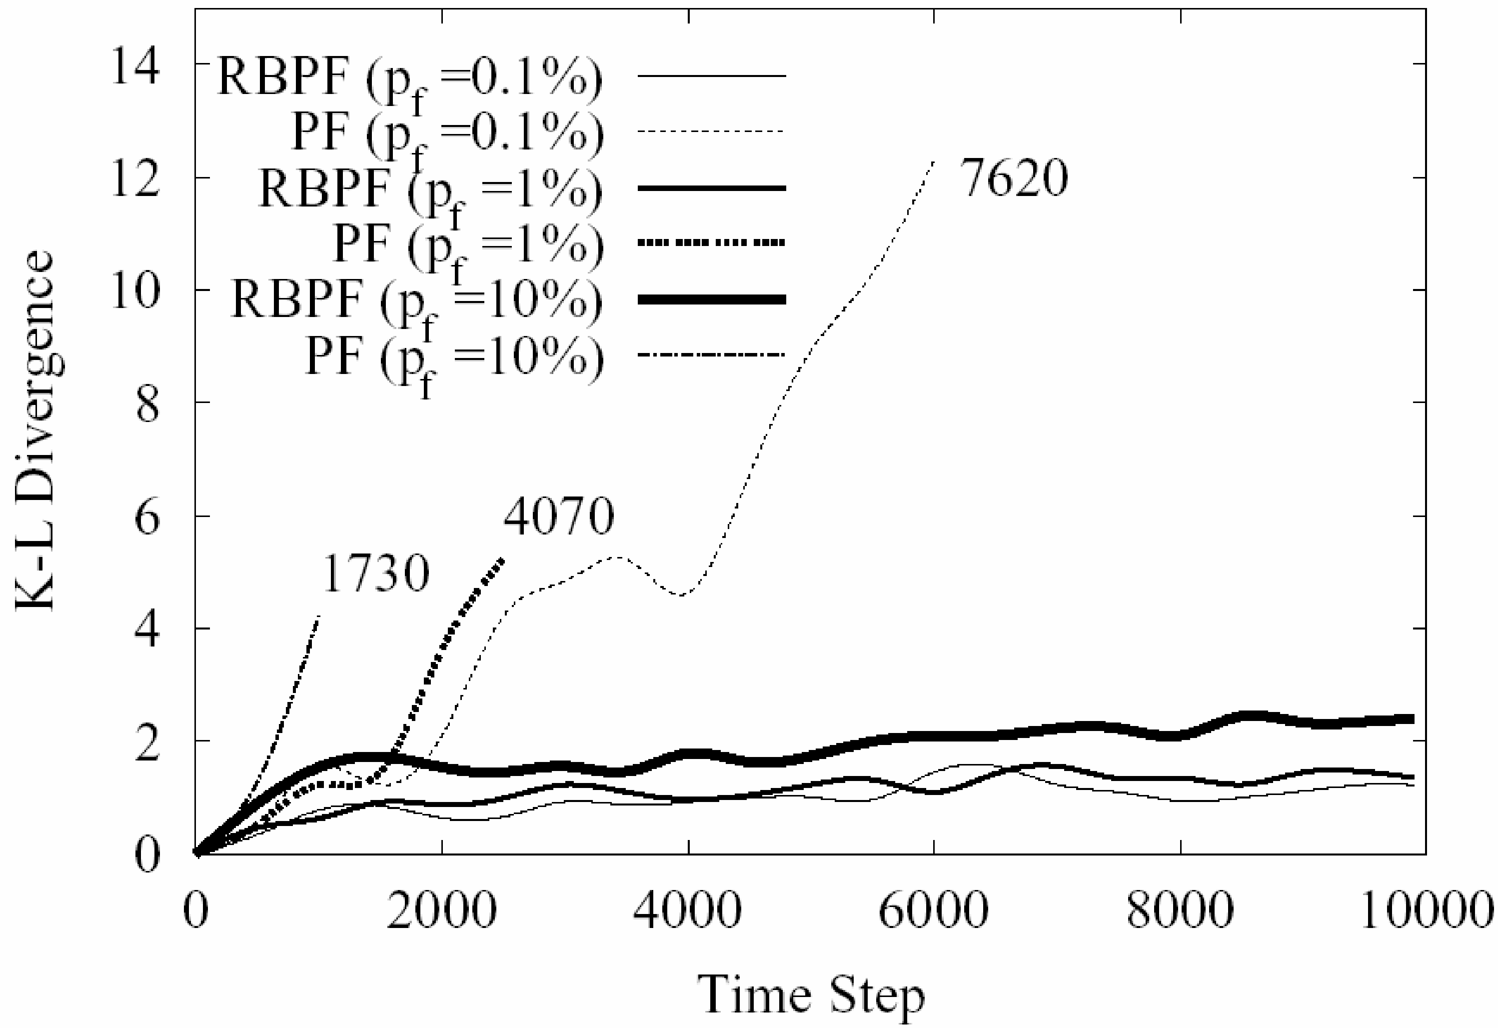

Figure 1: Comparison of RBPF (5000 particles) and PF (200,000 particles) for 1000 objects and varying fault probability.

Results

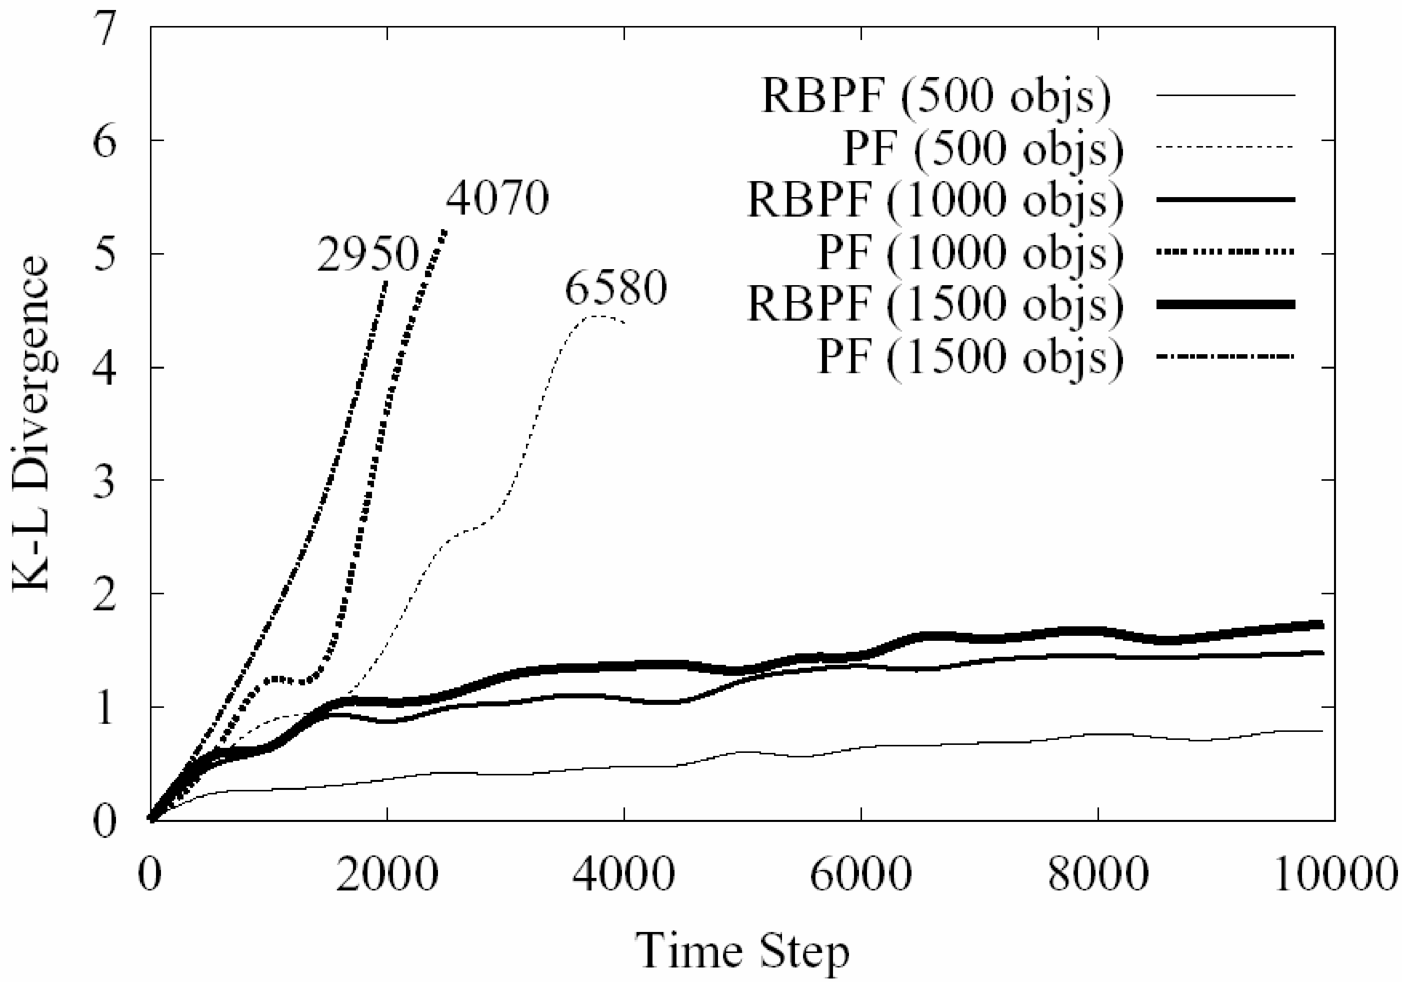

Figure 2: Comparison of RBPF (5000 particles) and PF (200,000 particles) for fault probability of 1% and varying number of objects.

Conclusion

Dynamic Probabilistic Relational Models (DPRMs), can represent uncertainty, sequential phenomena and relational structures.

With some strong assumption over the model, a scalable inference procedure based on Rao-Blackwellisation is provided.